



# DIABETIC RETINOPATHY SYMPTOMS, DIAGNOSE & TREATMENT

**Mrs. V. Tejaswini**

Dept. of CSE (Data Science) CMR Technical  
Campus Telangana, Hyderabad, India

**P. Sravanthi**

Dept. of CSE (Cyber Security)  
CMR Technical Campus  
Telangana, Hyderabad, India

**T. Omkar Babu**

Dept. of CSE (Cyber Security) CMR  
Technical Campus Telangana, Hyderabad,  
India

**V. Abhi Ram**

Dept. of CSE (Cyber Security) CMR  
Technical Campus Telangana, Hyderabad,  
India

**Abstract**— This layer turns the light that the lens focuses into signals that the brain can understand. The macula sits right in the center of your retina and handles the heavy lifting when it comes to sensing light. It picks up details, passes that info to the rest of the retina, and then the retina sends everything up the optic nerve to your brain—basically, that’s how you see. But when someone has age-related macular degeneration (AMD), optic disc drusen, Rothspot, diabetic macular edema (DME). The hardest part is finding eye diseases because they are so different that they can only be accurately diagnosed by an optometrist with a lot of experience. In the same way, eye diseases are easy to spot and treat early on with computer-aided diagnosis tools (CAD). Recently, a lot of cutting-edge machine learning (ML) and deep learning (DL) models have been put forward to help classify, segment, and identify eye illnesses. We see that identifying and collecting data are big problems when using ADDs.

**Keywords**—programming education, CodeMentor AI, interactive learning, real-time compilation, AI tutoring

## 1. INTRODUCTION

Retinal diseases are shows more and more, and it’s not just older adults people of all ages are dealing with them. This shows what’s happening inside the eye. The retina is a thin layer of nerve tissue at the back, picks up light focused by the lens and turns it into signals the brain understands. Right in the center, you’ve got the macula. It handles the actual sensing, gathers the information, and then passes it along through the optic nerve so you can see the world around you. Now, when something goes wrong like age-related macular degeneration (AMD), optic disc drusen, Rothspot, or diabetic macular edema (DME) vision takes a hit. The tough part? Spotting these diseases isn’t easy. They look so different from one another that only a skilled optometrist can really tell them what the problem is. But technology is starting to help. Computer-aided diagnosis (CAD) tools are making it easier to catch eye diseases early, and now there’s a wave of new machine learning (ML) and deep learning (DL) models designed to classify and identify these issues.

The most difficult part is catching these diseases early. They can look nearly different, so you really need an expert eye to know about the differences. But now? Technology’s stepping in. Computer-aided diagnosis (CAD) tools are picking up the slack, and there’s a whole new wave of machine learning (ML) and deep learning (DL) models out there, built to help doctors classify and identify eye diseases faster and with more accuracy. There’s still a big hurdle: actually getting the right data and pinning down the correct diagnosis. There’s no shortage of fancy models—Recurrent Neural Networks (RNNs), Convolutional Neural Networks (CNNs), AlexNet, ResNet, VGG—these have all moved the needle for researchers and doctors tackling these tough cases. That’s where this project comes in. The plan? Use advanced neural networks to crack the toughest parts of classifying serious eye diseases. Sure, things like U-Net Segmentation have made a difference, but they still slip up—they’re not great at pulling out features, and they eat up way too much memory and computing power. So, this project is rolling out a new CNN model, one that’s built for multi-class classification but uses memory way more efficiently. The main goal is to put this new model to the test using Eye Net, a big dataset that covers 32 different eye diseases. The real challenge is to show that this new approach accurate classifications and manages memory better than the older models. To prove the project looks at precision, overall accuracy, how the model work with different numbers of training epochs, and how fast it can process each step. Make deep learning a reliable and fast way to diagnose eye diseases. Medical imaging is a plan an important role when it comes to catching diseases early and figuring out treatment. That’s especially true for eye care—getting quick, accurate answers can change everything for a patient. The old way meant doctors staring at images for hours, which took forever and left too much room for human error. Now, with AI and deep learning, everything’s speeding up. Neural networks can scan medical images in a flash.

We are working on this because we need to fix these problems and make a better CNN model for classifying eye diseases into the different groups. We want to improve the accuracy and flexibility of disease identification systems by making better use of memory and computing power. Our suggested CNN model is better in a number of ways than current methods. By carefully planning how to best use memory, we not only lower the amount of memory needed, but we also boost performance in areas like precision, remember, and accuracy. Also, our model works well with a variety of epoch numbers, so the performance stays the same over time. We tested our suggested method on Eye-net, a common dataset that includes 32 types of eye diseases. This shows that it works in real life. By testing our project method against other old methods, we show that it works better than them, which supports the idea that our approach could completely change how eye diseases are classified. In the end, our work tackles some of the biggest challenges in medical image analysis. We speed up disease classification and make it more accurate. By getting the most out of CNNs and using memory efficiently, we set up smarter clinical decision support systems. These tools help doctors diagnose faster and more accurately, which leads to better patient outcomes and better care.

## 2. LITERATURE REVIEW

We dug deep into the latest research on how people verify and create certificates. Before putting this all literature together, we pored over all kinds of sources research papers, journals, even different media outlooks. Getting the labeled data in medical imaging is very important. Because it's a crucial step for both diagnosis and treatment. Deep learning has made a big difference here, too. It's done an impressive job with tasks like spotting skin cancer, diabetic retinopathy, age-related macular degeneration (AMD), and even detecting COVID-19 from chest X-rays. This study looks at the latest progress in using deep learning to classify medical images. It does this by looking at the methods, datasets, and performance measures that have been described in the literature. This is because medical picture analysis is a key part of helping doctors diagnose diseases quickly and correctly. Deep learning, especially convolutional neural networks (CNNs), has become a strong tool for medical picture analysis that can be used on a large scale and with high accuracy. Deep Learning Models for Skin Cancer Classification: Esteva et al. suggested using deep neural networks to classify skin cancer at the level of a physician. They had great success telling the difference between different types of skin tumors.

The literature shows a lot of information, but there are still some big disadvantages. One of the biggest issues? CNN models eat up a ton of memory and processing power. Take U-Net and ResNet, for example. They're powerful, sure, but they need a lot of resources just to handle high-resolution medical images. U-Net is especially demanding—it passes full feature maps between its encoder and decoder, which is a memory hog. In real clinical settings, that kind of inefficiency just doesn't work. There is another problem which is not cleared by old method. these systems don't always work's well. Small training datasets, class imbalances, and different images are captured all make it harder for models to perform consistently. So, there's a real need for smarter model's ones that keep accuracy high,

but don't waste memory or processing power. Lately, researchers have started looking at hybrid and lightweight models. MobileNet, Xception, and EfficientNet all do more with less information. For example, MobileNet uses depthwise separable convolutions to cut down on the number of computational cost without losing accuracy. Xception, which builds the Inception architecture, goes even further with depthwise convolutions, and it's made a mark in medical imaging. Some teams are even combining features from different lightweight models, and they're seeing better robustness and precision—especially for things like retinal disease detection.

### 2.1 EXISTING SYSTEM

Labeling medical images really matters for both diagnosis and treatment, so it's got to be done right. Deep learning has actually changed the game here. It's helped with all sorts of tasks—like spotting skin cancer, catching diabetic retinopathy, diagnosing age-related macular degeneration, and even detecting COVID-19 from chest X-rays. These methods have pulled off some pretty impressive results. This study looks at the latest progress in using deep learning to classify medical images. It does this by looking at the methods, datasets, and performance measures that have been described in the literature. This is because medical picture analysis is a key part of helping doctors diagnose diseases quickly and correctly.

Deep learning, especially convolutional neural networks (CNNs), has become a strong tool for medical picture analysis that can be used on a large scale and with high accuracy. Deep Learning Models for Skin Cancer Classification: Esteva et al.

suggested using deep neural networks to classify skin cancer at the level of a physician. They had great success telling the difference between different types of skin tumors. Classification of Diabetic Retinopathy: Shankar et al. created an automated system using a synergic deep learning model to find and categorize diabetic retinopathy.

This shows that it could be useful for early detection and planning treatment. Farsiou and his team came up with a mathematical approach that uses optical coherence tomography to spot the difference between eyes with and without intermediate age-related macular degeneration. On the other hand, Guefrechi and colleagues turned to deep learning to detect COVID-19 in chest X-rays. Their work showed that these AI methods really help healthcare systems, especially when the pandemic hit hard.

### 2.2 PROPOSED SYSTEM

We're introducing a convolutional neural network (CNN) model that handles multi-class problems and uses memory efficiently. To test it, we used Eye Net—a standard dataset with 32 different eye diseases. This deep learning-based CNN aims to improve how doctors usually diagnose these conditions, which really matter. Our model works a lot like the latest state-of-the-art approaches out there. We compared everything using precision, memory usage, and accuracy, tracking how long each step took and how many epochs we needed.

### 3. PROBLEM STATEMENT AND OBJECTIVES

Retinal diseases like diabetic retinopathy, glaucoma, and age-related macular degeneration can rob people of their sight if doctors don't catch them early. The problem is, diagnosing these issues by examining retinal images takes a lot of time. It also depends heavily on the skill and experience of the ophthalmologist looking at the scans. Even then, mistakes happen, and there just aren't enough specialists to keep up with the number of patients needing regular screening.

Looking at these images isn't easy, either. Retinal scans are full of tiny details and subtle signs that can easily slip by, especially in the early stages. Classic image-processing methods just aren't good enough for this—they miss things. Automated systems built on deep learning sound promising, but right now, most of them are bulky. They eat up a lot of memory, slow down computers, and often stumble when asked to spot several different diseases at once, especially if those diseases look similar.

That's where this project comes in. The goal is to build a fast, lightweight, and accurate deep learning system called CNN model. That can reliably detect up to 32 different retinal diseases from the EyeNet dataset without bogging down computers or needing lots of memory. If this technique is used, doctors can diagnose patients more quickly, especially in places where specialists and resources are stretched. It'll speed up screenings, reduces down time and mistakes, and help make sure people get the treatment they need before their vision worsens.

Our aiming is:

- Build a deep CNN that can spot and accurately classify retinal diseases from fundus images.
- Gather and clean up a big dataset of retinal images for training and testing.
- Use the CNN to pull out meaningful features and patterns linked to specific diseases.
- Automatically sort images into "healthy" or "diseased" groups, so doctors get fast, reliable answers.
- Lighten the workload for ophthalmologists, so they can see more patients and catch problems earlier.
- Measure how well the system works using clear metrics like accuracy, precision, recall, and F1-score, making sure it's actually useful in real-world clinics.

This project isn't just about building a clever algorithm. It's about making early detection of retinal disease faster, more reliable, and accessible. So fewer people have to face preventable vision loss.

- Gather and preprocess retinal image datasets to boost how well the model learns and gets tested.
- Use CNN layers to pull out important features and spot disease patterns.
- Automatically sort retinal images into "diseased" or "healthy" groups.
- Cut down on manual work so ophthalmologists can diagnose patients faster.
- Catch diseases earlier and help prevent long-term vision loss by stepping in sooner.
- Check how the model performs by looking at accuracy, precision, recall, and the F1-score.

### 4. SYSTEM ARCHITECTURE AND DESIGN

The suggested system structure uses a convolutional neural network (CNN) model that is best for sorting eye diseases into multiple groups. Innovative memory management methods are built into the design to improve memory efficiency. The design is made up of separate parts that work together to make feature extraction and classification work well. Unlike other designs, this one reduces the amount of extra memory needed by making the flow of data more efficient. An analysis of the Eye Net dataset, which includes 32 types of eye diseases, shows that it performs better in terms of accuracy and resource use. Key measures like memory, precision, and total accuracy show that the suggested method works.

This system takes a pretty practical approach to spotting and sorting out diseases in medical images. It starts by collecting retinal images, then cleans them up by resizing, normalizing, removing of noise, and mixing in some data augmentation to make sure everything's sharp and consistent. Once this process is done, it feeds the images into a deep Convolutional Neural Network. The several convolution and pooling layers dig out key features, like lesions or anything abnormal. After that, the system hands these features off to fully connected layers and uses a Softmax or Sigmoid classifier to nail down which disease category fits. The training process is thorough—lots of testing and tweaking to get the model as accurate and reliable as possible. When the model is ready, the model can also look at new retinal images and predict diseases fast. If there is a need, you can also add a simple user interface so people can upload images and see results right away. Finally, you can run the trained model on the cloud or on local servers, making medical screening both efficient and accurate.

We train and test the system carefully to make our model accurate and reliable. Then, the model can look at new retinal images and predict diseases. There is also an option for a user interface, so that the people can upload images and see the results easily. When it is ready, you can set up the trained model on the cloud or on your own servers, which makes medical screening faster and more accurate. Retina disease detection systems use deep learning to identify eye conditions like diabetic retinopathy, glaucoma, and macular degeneration from retinal fundus images. It starts when someone uploads an image or snaps one with a fundus camera. The system stores these images securely, then runs them through several steps called resizing, noise reduction, boosting contrast, normalization, and sometimes data augmentation to keep the model sharp.

Retina detection systems use deep learning to know eye conditions like diabetic retinopathy, glaucoma, and macular degeneration from retinal fundus images. It starts when someone uploads an image or snaps one with a fundus camera. The system stores these images securely, and runs them through several steps: resizing, noise reduction, boosting contrast, normalization, and sometimes data augmentation to keep the model sharp. Once it trained, the model sits behind an inference service. Basically, it's ready to handle real-time predictions through APIs, so results come back fast and reliably. To help doctors trust what's happening, the system uses tools like Grad-CAM to create heatmaps that show which parts of the retina look suspicious. All the results predictions, confidence scores, and those visual explanations get pulled together in a web dashboard or app for clinicians.

This system is built to scale, stay secure, and follow the rules. It helps an eye on performance, retrains itself with new data and fits right into large-scale screenings or tele-ophthalmology work.

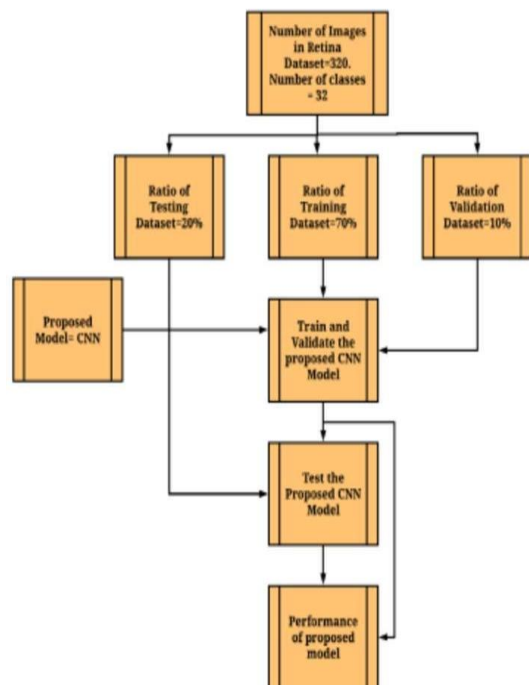


Figure 1: CNN Training, Testing.

#### A. DATA FLOW

##### 1. Input Image

- The process begins with a raw **retinal fundus image** captured from an eye-screening device.
- This image may contain variations in lighting, clarity, and noise, which need improvement before analysis.

##### 2. Image Pre-processing

- The raw image undergoes enhancement techniques to improve its quality.
- Steps may include:
  - Noise removal
  - Contrast improvement
  - Color normalization
  - Resizing
- The output is a **cleaner, enhanced retinal image** suitable for feature extraction.
- These features are essential indicators of diseases like Diabetic Retinopathy or Glaucoma.

**YOLO v7 Feature Extraction** First, the enhanced retinal image goes straight into the YOLO v7 model. This algorithm doesn't just look for anything — it hunts for key retinal features: microaneurysms, exudates, hemorrhages, vessel structures, and optic disc patterns. These details matter. They're the signs where doctors look for when diagnosing diseases like Diabetic Retinopathy or Glaucoma.

Once YOLO v7 pulls out those features, they all get bundled together into a feature set. But not every feature plays a useful

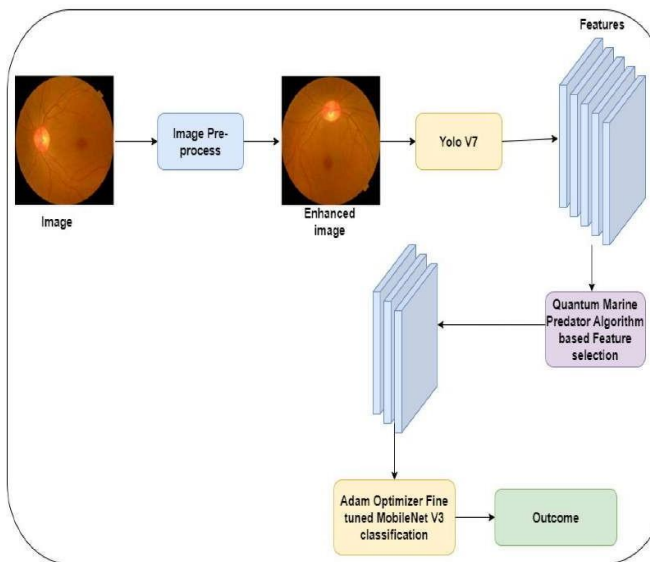
role. Some just add clutter, so the next step is trimming the fat.

**Quantum Marine Predator Algorithm – Feature Selection** Here's where things get smart. The Quantum Marine Predator Algorithm (QMPA) steps in to pick out only the features that actually help with diagnosis. This step cuts out noise, ditches irrelevant info, and keeps just the essentials. The pay-off? The system runs faster and gets more accurate results. After this, the optimized feature set heads to the classifier. Adam Optimizer + Fine-Tuned MobileNet V3 Classification Now, the refined features feed into a specially fine-tuned MobileNet V3 model. MobileNet V3 is fast, lightweight, and built for jobs like medical imaging. The Adam optimizer helps the model learn quicker, converge faster, and hit higher accuracy. The classifier takes the cleaned features and makes its call. Outcome is, the system gives you a clear result healthy, Diabetic Retinopathy, Glaucoma, Age-related Macular Degeneration (AMD), or maybe another retinal disorder.

This outcome comes from the combined effort of image enhancement, careful feature extraction, smart optimization, and deep learning classification. In short, the proposed system for retinal disease detection follows a sharp, step-by-step pipeline. It starts by cleaning up the raw images, then uses YOLO v7 to zero in on the key retinal features. Next, QMPA filters out what doesn't matter, leaving just what's important. Those optimized features go into a fine-tuned MobileNet V3 model, powered by Adam for better learning. The result? Efficient, accurate, and reliable detection so you know exactly what's going on with the retina, every time. Detecting retinal diseases with deep convolutional neural networks—CNNs, it starts with a simple idea, they teach a computer to identify eye problems just by looking at retinal photos.

First, take image of the retina with a fundus camera. But raw images can make the system damage, so you clean them up. Resize, normalize, clear out the noise, bump up the contrast these steps make important details like blood vessels or tiny lesions stand out. Then, the magic happens inside the CNN. At first, the network will take a basic stuff like edges, textures, little shapes. As you move deeper into the layers, it starts to notice more complex patterns. Things like microaneurysms, hemorrhages, or exudates.

These are big clues for identifying different retinal diseases. The network keeps refining what it sees, layer by layer, through a mix of operations—convolution, activation, pooling. Eventually, it boils everything down and pushes the info into fully connected layers that make the final call: disease or no disease, maybe even how severe it is. Training the CNN is a grind. You must keep a bunch of labeled images, and it keeps adjusting its settings, trying to get better at matching the right diagnosis. It learns by minimizing mistakes, this all done by backpropagation and a loss function that tells it when it's wrong. Once the model's ready, you can throw new retinal images at it. It'll spit out predictions—and not just a yes or no, but confidence scores too. If you want to see what the model's thinking, you can use tools like Grad-CAM, which actually highlights the troubled spots in the image. That helps doctors understand and trust what the AI is seeing, making the diagnosis a bit less of a black box. But the work doesn't stop after deployment. Teams keep an eye on the model's predictions to catch any drop in performance, which sometimes happens if image quality changes or the patient population shifts.



**Figure 2: Overall System Architecture for Retinal Disease Detection**

The proposed system for retinal disease detection uses a structured deep-learning pipeline to accurately analyze retinal fundus images which are collected by dataset. The process begins with image preprocessing, where raw retinal images are enhanced to improve clarity and remove noise data from images.

## 5. IMPLEMENTATION

### FRONT-END IMPLEMENTATION

#### 1. User Interface (UI) Design

Create a simple and user-friendly interface where users can:

- Upload a retinal fundus image
- View preprocessing output
- View detection results
- Download or save the final diagnosis report

#### UI Framework Options:

- **Web:** HTML, CSS, JavaScript, Bootstrap, React
- **Desktop:** Tkinter / PyQt / Streamlit

#### 2. Image Upload Module

The front end must include:

- A button to choose image from device
- A preview of the uploaded image
- Validation to ensure only image formats like .jpg, png are accepted

#### Example Feature:

- “Drag and Drop to Upload Retina Image”

### BACKEND IMPLEMENTATION

#### 1. Backend Environment Setup

Prepare the backend environment with required libraries:

##### Main Libraries

- Python
- OpenCV

- NumPy
- TensorFlow / PyTorch
- YOLOv7 (PyTorch)
- MobileNet V3
- Scikit-learn
- Image processing libraries
- Flask / FastAPI (to connect to the frontend)

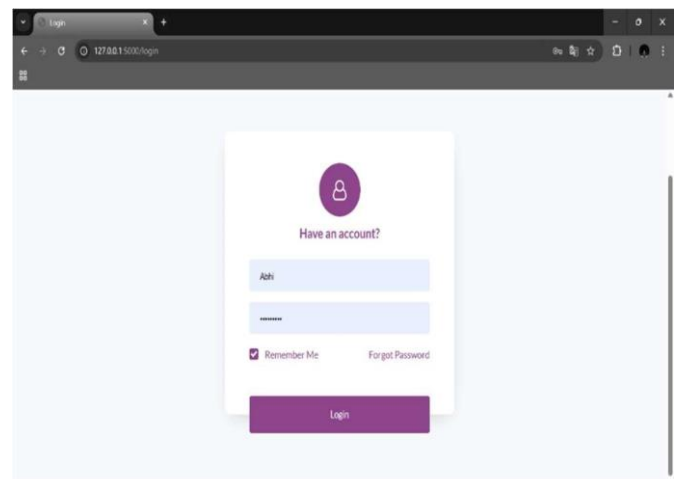
#### 2. Dataset Handling Module

The backend must:

- Load the retinal image dataset
- Split images into **train, test, validation**
- Store file paths
- Provide images to the CNN models for training/testing  
This is automated through Python scripts.

## 6. RESULTS AND DISCUSSION

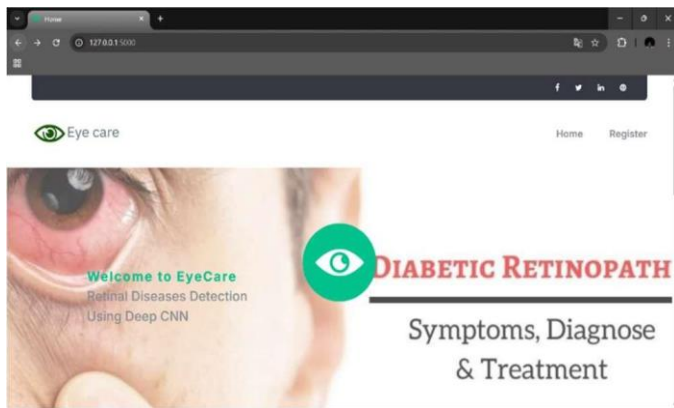
The study presents a new CNN model that is intended to efficiently use memory for multi-class classification of eye diseases. Looking at the EyeNet collection, which has 32 eye diseases, shows encouraging results. Compared to current state-of-the-art methods, the suggested model does a better job of managing memory and being accurate. Researchers did some tests and found that the model had good validation accuracy across a range of epochs. It did especially well at 10 and 15 epochs, with only a small change in validation loss (0.0279) between them.



**Figure 3: Login page of the Retinal detection system displaying**

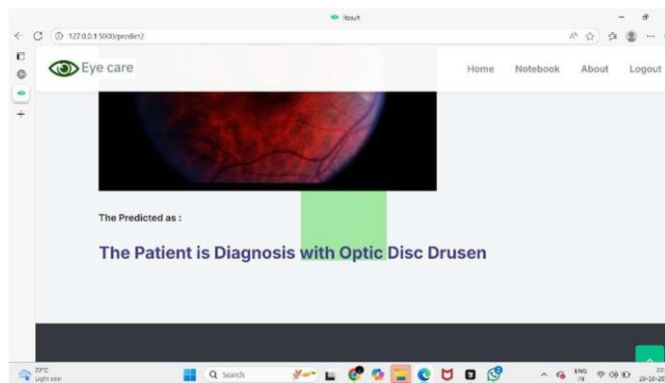
The image shows the homepage of a retinal detection web application. Which focused on retinal disease detection using Deep CNN. It show user login page. Where the user can login with their name and enter the application page.

If user name and password is correct, you can login into the page.



**Figure 4: Homepage of the Retinal detection system showing its interface for retinal disease detection**

The image shows the homepage of a retinal detection system web application focused on retinal disease detection using Deep CNN. The page shows a large banner image of an eye on the left and text about Diabetic Retinopathy on the right. The top navigation bar includes options like Home and Register, along with social media icons.



**Figure 5: Result Page Displaying Model Prediction Optic Disc Drusen Diagnosis**

The image shows the results page of the EyeCare web application. It displays a retinal scan image at the top, followed by a prediction message stating that the patient is diagnosed with Optic Disc Drusen. The page includes a navigation bar with options like Home, Notebook, About, and Logout.

## 7. CONCLUSION AND FUTURE WORK

We show a CNN model based on deep learning that can sort the different eye diseases into groups. EyeNet, a collection with 32 different eye diseases, is used to put the model into action. Different epochs are used to train the proposed model and test how well it works. At first, the model was trained over 10 iterations and had good validation accuracy. Then, it was trained over 15 iterations and again had good validation accuracy with a validation loss of 0.0279, which is different each time. The model's overall performance is much better than that of other models that are thought to be the best on the market. That being said, the model that was given might help with the classification of eye illnesses. Updating the model often and retraining it with new data will keep making it work better in the future. This will be done by taking advantage of improvements in deep learning methods and the growing number of datasets that include different types of eye diseases.

## FUTURE ENHANCEMENTS AND DISCUSSIONS

There are a number of ways that the CNN model for classifying eye diseases can be improved in the future. To begin, using transfer learning methods on bigger and more varied picture datasets can improve performance even more by using models that have already been taught. By mixing results from different models, ensemble learning methods could also improve the stability and generalizability of models. Adding self-attention or attention processes may also help the model focus on important parts of retinal pictures, which could lead to better classification accuracy. It is important to keep improving models by fine-tuning them on new and bigger datasets so that they can adapt to changing trends in eye diseases. Using new deep learning structures and methods, like graph neural networks or capsule networks, could help us learn new things and make our classifications more accurate. Lastly, putting the model to use in hospital settings and getting feedback from doctors and nurses can help make more changes that fit the needs of real-world applications.

## 8. REFERENCES

- [1] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H.M. Blau, and S. Thrun, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, Feb. 2017.
- [2] K. Shankar, A. R. W. Sait, D. Gupta, S. K. Lakshmanaprabu, A. Khanna, and H. M. Pandey, "Automated detection and classification of fundus diabetic retinopathy images using synergic deep learning model," *Pattern Recognit. Lett.*, vol. 133, pp. 210–216, May 2020.
- [3] R. Arunkumar and P. Karthigaikumar, "Multi-retinal disease classification by reduced deep learning features," *Neural Comput. Appl.*, vol. 28, no. 2, pp. 329–334, Feb. 2017.
- [4] T. Shanthy and R. S. Sabeenian, "Modified Alexnet architecture for classification of diabetic retinopathy images," *Comput. Electr. Eng.*, vol. 76, pp. 56–64, Jun. 2019.
- [5] S. Farsi, S. J. Chiu, R. V. O'Connell, F. A. Folgar, E. Yuan, J. A. Izatt, and C. A. Toth, "Quantitative classification of eyes with and without intermediate age-related macular degeneration using optical coherence tomography," *Ophthalmology*, vol. 121, no. 1, pp. 162–172, Jan. 2014.
- [6] R. F. Mullins, S. R. Russell, D. H. Anderson, and G. S. Hageman, "Drusen associated with aging and age-related macular degeneration contain proteins common to extracellular deposits associated with atherosclerosis, elastosis, amyloidosis, and dense deposit disease," *FASEB J.*, vol. 14, no. 7, pp. 835–846, May 2000.
- [7] Y. Kanagasingam, A. Bhuiyan, M. D. Abramoff, R. T. Smith, L. Goldschmidt, and T. Y. Wong, "Progress on retinal image analysis for age related macular degeneration," *Prog. Retinal Eye Res.*, vol. 38, pp. 20–42, Jan. 2014.
- [8] D. S. Kermay, "Identifying medical diagnoses and treatable diseases by image-based deep learning," *Cell*, vol. 172, no. 5, pp. 1122–1131, Feb. 2018.
- [9] M. M. M. S. Fathy and M. T. Mahmoudi, "A classified and comparative study of edge detection algorithms," in *Proc. Int. Conf. Inf. Technol., Coding Comput.*, Apr. 2002, pp. 117–120.
- [10] C.-H. H. Yang, J.-H. Huang, F. Liu, F.-Y. Chiu, M. Gao, W. Lyu, M. D. I.-H. Lin, and J. Tegner, "A novel hybrid machine learning model for auto-classification of retinal diseases," 2018, arXiv:1806.06423.